



Can AI fix production incidents?

60+ open-source contributors | <https://sregym.com/> | NeurIPS 2026 paper

Overview and Value Proposition

A live RL environment for AI SRE agents

- Agents diagnose and fix injected faults in live Kubernetes stacks

A benchmark frontier agents have not saturated

- Frontier agents resolve 52–81% of SREGym-Lite incidents end to end



SREGym Problems

- 100+ live SRE problems

Fault	Layer	App
Webhook outage	Controller	Hotel
Bad readiness probe	Config	Social
Service DNS failure	Network	Social
Retry storm collapse	Metastable	Hotel
Stale rotated secret	Config	Astronomy

* 5 of the 21 SREGym-Lite faults shown
* Also: OS, hardware, concurrent, metastable faults

SREGym for RL

Verifiable rewards

- Oracles grade diagnosis and mitigation separately
- Graders hidden from the agent
- Private cluster per rollout

Scalable task set

- Replays real outages: Cloudflare, Reddit, GKE, Kafka

SREGym 2.0: Environment Scaling

Coding agents build executable copies of real production systems at several scale tiers

Realistic, long-horizon, difficult

- Realistic:** each task rebuilds a real postmortem with its failure and recovery logic
- Long-horizon:** the incident keeps evolving, with a recovery tail after the first fix
- Difficult:** tiers add size, dependency depth, noisy telemetry, broken tools, ambiguity

First milestone: 10 validated tasks across 4–6 incident families

First candidate postmortems

- Roblox, 2021:** Consul contention took the platform down for **73 hours**
- Datadog, 2023:** a systemd update wiped Cilium routes across regions
- Rogers, 2022:** BGP routes leaked into OSPF, downing the core and management plane

Results

SREGym-Lite-0904 leaderboard, 21 faults (sregym.com, Oct 2026)
End-to-end = diagnosis and mitigation both correct in the same run

Model (Agent)	End-to-end	Diagnosis	Mitigation	Time to fix	Tokens per run
 GPT-5.6 Sol, max (Codex)	81.0%	95.2%	85.7%	6.6 min	1.42M
 Claude Opus 5 (Claude Code)	76.2%	92.1%	82.5%	7.0 min	1.64M
 Claude Opus 4.8 (Claude Code)	52.4%	66.7%	63.5%	8.4 min	1.66M

- Scores track progress:** +24 points from Claude Opus 4.8 to Opus 5
- Short episodes:** top agents fix an incident in about 7 minutes

Open-source community 60+ contributors 300+ GitHub stars 100+ forks